

WE CLAIM:

1. A method of dynamically switching among a plurality of fault tolerance schemes associated with a fault tolerance mechanism that executes in a distributed system, the method comprising:

5 obtaining a wait time of at least one user interface event occurring in said distributed system, said wait time including at least one of a communications time, a service time and a fault tolerance time;

determining whether a mean of said wait time is greater than a predetermined mean wait time threshold;

10 determining whether said communications time, said service time and said fault tolerance time are mutually independent when said mean of said wait time is greater than said predetermined mean wait time threshold;

determining whether said mean of said wait time can be improved by reducing a mean of said fault tolerance time when said communications time, said service time and said fault tolerance time are mutually independent; and

15 switching from a first of said plurality of fault tolerance schemes to a second of said plurality of fault tolerance schemes when said wait time can be improved by reducing said mean of said fault tolerance time.

20 2. The method of claim 1 wherein said mean wait time threshold is set by an application associated with said at least one user interface event.

3. The method of claim 2 wherein said application specifies said mean wait time threshold per class of user interface events associated with said application.

25 4. The method of claim 2 wherein said mean wait time threshold set by said application can be changed by a user of said application.

5. The method of claim 1 wherein said mean wait time threshold is set using a profile of a user of an application associated with said at least one user interface event.

6. The method of claim 5 wherein said mean wait time threshold is set using said user profile on a per device basis.

5 7. The method of claim 1 wherein said determining whether said communications time, said service time and said fault tolerance time are mutually independent includes determining whether a number of messages being passed between devices executing an application associated with said at least one user interface event exceeds a predetermined message threshold.

10 8. The method of claim 1 wherein said determining whether said mean of said wait time can be improved by reducing a mean of said fault tolerance time includes determining whether said mean of said fault tolerance time as a percentage of said mean of said wait time is greater than a predetermined fault tolerance threshold.

15 9. The method of claim 1 wherein said switching from a first of said plurality of fault tolerance schemes to a second of said plurality of fault tolerance schemes further includes selecting said second fault tolerance scheme based on a set of requirements specified by an application associated with said at least one user interface event.

20 10. The method of claim 1 wherein said switching from a first of said plurality of fault tolerance schemes to a second of said plurality of fault tolerance schemes further includes selecting said second fault tolerance scheme based on implementation costs associated with at least one of said first and second fault tolerance schemes.

25 11. The method of claim 1 further comprising:  
determining a cost for said first fault tolerance scheme when said communications time, said service time and said fault tolerance time are not mutually independent;  
determining whether said first fault tolerance scheme has a significant impact on said communication time and said service time; and

switching from said first fault tolerance schemes to a third of said plurality of fault tolerance schemes when said significant impact is determined.

5           12.    The method of claim 11 wherein said significant impact is defined as an effect on said communications time and service time that causes said mean wait time to increase above said mean wait time threshold.

10           13.    The method of claim 11 further comprising:  
              determining whether an application associated with said at least one user interface event requires an increased level of fault tolerance when said mean of said wait time is not greater than a predetermined mean wait time threshold;

              determining whether at least one fault tolerance scheme having said increased level of fault tolerance can meet said mean wait time threshold; and

15           switching from said first fault tolerance schemes to select one of said at least one fault tolerance scheme having said increased level of fault tolerance.

20           14.    The method of claim 13 further comprising notifying an application associated with said at least one user interface event of a decision to switch fault tolerance schemes.

              15.    A fault tolerant distributed system capable of dynamically switching among a plurality of fault tolerance schemes associated with a fault tolerance mechanism, the system comprising:

25           means for obtaining a wait time of at least one user interface event occurring in said distributed system, said wait time including a communications time, a service time and a fault tolerance time;

              means for determining whether a mean of said wait time is greater than a predetermined mean wait time threshold;

means for determining whether said communications time, said service time and said fault tolerance time are mutually independent when said mean of said wait time is greater than said predetermined mean wait time threshold;

5 means for determining whether said mean of said wait time can be improved by reducing a mean of said fault tolerance time when said communications time, said service time and said fault tolerance time are mutually independent; and

10 means for switching from a first of said plurality of fault tolerance schemes to a second of said plurality of fault tolerance schemes when said wait time can be improved by reducing said mean of said fault tolerance time.

16. The system of claim 15 further comprising:

15 means for determining a cost for said first fault tolerance scheme when said communications time, said service time and said fault tolerance time are not mutually independent;

means for determining whether said first fault tolerance scheme has a significant impact on said communication time and said service time; and

20 means for switching from said first fault tolerance schemes to a third of said plurality of fault tolerance schemes when said significant impact is determined.

17. The system of claim 16 further comprising:

25 means for determining whether an application associated with said at least one user interface event requires an increased level of fault tolerance when said mean of said wait time is not greater than a predetermined mean wait time threshold;

means for determining whether at least one fault tolerance scheme having said increased level of fault tolerance can meet said mean wait time threshold; and

means for switching from said first fault tolerance schemes to select one of said at least one fault tolerance scheme having said increased level of fault tolerance.

- 5           18.    The system of claim 17 further comprising means for notifying an application associated with said at least one user interface event of a decision to switch fault tolerance schemes.